APPENDIX

9. A computing system comprising:

a compiler for forming groups of instructions having opcodes including a first group of instructions and a second group of instructions, instructions in the first group of instructions executable in parallel, and instructions in the second group of instructions executable in parallel;

a first memory storage having at least a memory location, the memory location for storing the first group of instructions, for storing the second group of instructions comprising at least one instruction, and for storing group identifiers that indicate which instructions are included within the first group of instructions and which instructions are included within the second group of instructions;

a pre-decoder coupled to the first memory storage for decoding opcodes of instructions in the first group of instructions and opcodes of instructions in the second group of instructions, for forming a first group of expanded instructions, a second group of expanded instructions, and expanded group identifiers, and for determining processing pipeline identifiers associated with expanded instructions in the first group of expanded instructions and processing pipeline identifiers associated with expanded instructions in the second group of expanded instructions, in response thereto;

a second memory storage coupled to the predecoder having at least a memory location, the memory location for storing the first group of expanded instructions, the second group of expanded instructions, the expanded group identifiers, the processing pipeline identifiers associated with the expanded instructions in the first group of expanded instructions, and the processing pipeline identifiers associated with the expanded instructions in the second group of expanded instructions;

a decoder coupled to the second memory storage for receiving the first group of expanded instructions, the second group of expanded instructions, and the expanded group identifiers, and for issuing the first group of expanded instructions in response to the expanded group identifiers;

a plurality of processing pipelines;

a crossbar coupled to the decoder and to the plurality of processing pipelines, for issuing expanded instructions in the first group of expanded instructions to processing pipelines of the plurality of processing pipelines in response to the processing pipeline identifiers associated with the expanded instructions in the first group of expanded instructions.

10. The computing system of claim 9 wherein a first processing pipeline identifier from the processing pipeline identifiers associated with the expanded instructions in the first group of expanded instructions identifies a first floating point unit pipeline.

11. The computing system of claim 10 wherein a second processing pipeline identifier from the processing pipeline identifiers associated with the expanded instructions in the first group of expanded instructions identifies a first arithmetic logic unit pipeline.

12. The computing system of claim 10 wherein a second processing pipeline identifier from the processing pipeline identifiers associated with the expanded instructions in the first group of expanded instructions identifies a second floating point unit pipeline.

13. The computing system of claim 11 wherein a third processing pipeline identifier from the processing pipeline identifiers associated with the expanded instructions in the first group of expanded instructions identifies a second floating point unit pipeline.

14. The computing system of claim 9 wherein a first processing pipeline identifier from the processing pipeline identifiers associated with the expanded instructions in the second group of expanded instructions identifies a store unit pipeline.

15. The computing system of claim 14 wherein a second processing pipeline identifier from the processing pipeline identifiers associated with the expanded instructions in the second group of expanded instructions identifies a control unit pipeline.

16. The computing system of claim 9

wherein the decoder is also for issuing the second group of expanded instructions in response to the expanded group identifiers; and

wherein the crossbar is also for issuing expanded instructions in the second group of expanded instructions to processing pipelines of the plurality of processing pipelines in response to the processing pipeline identifiers associated with the expanded instructions in the second group of expanded instructions.

17. The computing system of claim 9 wherein the group identifiers are associated with instructions in the first group of instructions and with instructions in the second group of instruction.

18. The computing system of claim 17 wherein the group identifiers are embedded with instructions in the first group of instructions and with instructions in the second group of instruction.

19. The computing system of claim 9 wherein the first group of expanded instructions comprises at least one expanded instruction and the second group of expanded instructions comprises at least two expanded instructions.

20. The computing system of claim 19 wherein the second memory storage includes at least the two expanded instructions of the second group of expanded instructions and the one expanded instruction of the first group of expanded instructions.

21. The computing system of claim 9

wherein the compiler is also for forming a third group of instructions, the instructions in the third group of instructions executable in parallel;

wherein the memory location of the first memory storage is also for storing a third group of instructions in parallel with the first group of instructions and the second group of instructions; and

wherein the group identifiers indicate which instructions are included within the third group of instructions.

22. The computing system of claim 21 wherein the third group of instructions comprises at least one instruction.

23. The computing system of claim 9 wherein the first memory storage is a superscaler cache.

24. The computing system of claim 9 wherein one expanded instruction of the second group of expanded instructions is a branch instruction.

25. The computing system of claim 9 wherein the compiler explicitly identifies instructions that can be performed in parallel for the first group of instructions.

26. A method for issuing groups of individual software-scheduled instructions in parallel for processing comprises:

forming a first group of software-scheduled instructions, a second group of software-scheduled instructions comprising at least one instruction, and group identifiers indicating which software-scheduled instructions are included within the first group of software-scheduled instructions and which software-scheduled instructions are included within the second group of software-scheduled instructions, software-scheduled instructions in the first group of software-scheduled instructions having opcodes and executable in parallel, and software-scheduled instructions in the second group of software-scheduled instructions having opcodes and executable in parallel;

storing the first group of software-scheduled instructions, the second group of software-scheduled instructions, and the group identifiers in parallel in a first memory location;

forming a first group of expanded software-scheduled instructions, a second group of expanded software-scheduled instructions, and expanded group identifiers in response to opcodes of software-scheduled instructions in the first group of software-scheduled instructions and opcodes of software-scheduled instructions in the second group of software-scheduled instructions;

determining processing pipelines appropriate for expanded software-scheduled instructions in the first group of expanded software-scheduled instructions and processing pipelines appropriate for expanded software-scheduled instructions in the second group of expanded software-scheduled instructions also in response to the opcodes of software-scheduled instructions in the first group of software-scheduled instructions and the opcodes

of software-scheduled instructions in the second group of software-scheduled instructions; and

issuing the first group of expanded software-scheduled instructions to the processing pipelines appropriate for expanded software-scheduled instructions in the first group of expanded software-scheduled instructions, in response to the expanded group identifiers.

27. The method of claim 26 further comprising:

issuing the second group of expanded software-scheduled instructions to the processing pipelines appropriate for expanded software-scheduled instructions in the second group of expanded software-scheduled instructions, in response to the expanded group identifiers.

28. The method of claim 26 wherein a first processing pipeline appropriate for an expanded software-scheduled instruction in the first group of expanded software-scheduled instructions is coupled to a first arithmetic logic processing unit.

29. The method of claim 28 wherein a second processing pipeline appropriate for an expanded software-scheduled instruction in the first group of expanded software-scheduled instructions is coupled to a first floating point processing unit.

30. The method of claim 28 wherein a second processing pipeline appropriate for an expanded software-scheduled instruction in the first group of expanded software-scheduled instructions is coupled to a second arithmetic logic processing unit.

31. The method of claim 29 wherein a third processing pipeline appropriate for an expanded software-scheduled instruction in the first group of expanded software-scheduled instructions is coupled to a second arithmetic logic processing unit.

32. The method of claim 28 wherein a second processing pipeline appropriate for an expanded software-scheduled instruction in the first group of expanded software-scheduled instructions is coupled to a load unit.

33. The method of claim 26 wherein a first processing pipeline appropriate for an expanded software-scheduled instruction in the second group of expanded software-scheduled instructions is coupled to a store unit.

34. The method of claim 33 wherein a second processing pipeline appropriate for an expanded software-scheduled instruction in the second group of expanded software-scheduled instructions is coupled to a control unit.

35. The method of claim 26 further comprises:

storing the first group of expanded software-scheduled instructions, the second group of expanded software-scheduled instructions, and the expanded group identifiers, in a second memory location;

wherein the step of issuing the first group of expanded software-scheduled instructions comprises issuing the first group of expanded software-scheduled instructions and the second group of expanded software-scheduled instructions from the second memory location to a decoder; and

issuing the first group of expanded software-scheduled instructions to the processing pipelines appropriate for the expanded software-scheduled instructions in the first group of expanded software-scheduled instructions from the decoder in response to the expanded group identifiers stored in the second memory location.

36. The method of claim 35 further comprises:

issuing the second group of expanded software-scheduled instructions to the processing pipelines appropriate for the expanded software-scheduled instructions in the second group of expanded software-scheduled instructions from the decoder in response to the expanded group identifiers stored in the second memory location.

37. The method of claim 36 wherein the processing pipelines appropriate for the expanded software-scheduled instructions in the first group of expanded software-scheduled instructions are identified by processing pipeline identifiers.

38. The method of claim 37 wherein the step of storing the first group of expanded software-scheduled instructions further comprises storing the processing pipeline identifiers in the second memory location.

39. The method of claim 38 wherein the step of issuing the first group of expanded software-scheduled instructions to the processing pipelines comprises issuing the first group of expanded software-scheduled instructions to a crossbar.

40. The method of claim 39 wherein the crossbar is an associative crossbar responsive to the processing pipeline identifiers.

41. The method of claim 35 wherein the first group of expanded instructions comprises at least two expanded instructions.

42. The method of claim 41 wherein the second group of expanded instructions comprises at least one expanded instruction.

43. The method of claim 42 wherein the two expanded instructions of the first group of expanded instructions and the one expanded instruction of the second group of expanded instructions are stored in the second memory location

44. The method of claim 26 wherein the step of forming the first group

of software-scheduled instructions, the second group of software-scheduled instructions, and the group identifiers comprises using a compiler to form the first group of software-scheduled instructions, the second group of software-scheduled instructions, and the group identifiers.

45. The method of claim 44 wherein the group identifiers are associated with instructions in the first group of instructions and with instructions in the second group of instruction.

46. The method of claim 44 wherein the compiler explicitly determines parallel executable instructions among a plurality of instructions to include in the first group of software-scheduled instructions.

47. The method of claim 26 wherein the step of forming the first group of software-scheduled instructions, the second group of software-scheduled instructions, and the group identifiers further comprises the step of forming a third group of software-scheduled instructions executable in parallel; and

wherein the group identifiers also indicate which software-scheduled instructions are included within the third group of software-scheduled instructions.

48. The method of claim 26 wherein the first memory location is a cache location in a superscaler cache.

49. A method for issuing a group of individual instructions in parallel for processing comprises:

storing in parallel a plurality of instructions and instruction grouping information in a location in a memory, the plurality of instructions and the instruction grouping information determined by a compiler, the instruction grouping information indicating which instructions of the plurality of instructions belong to a first group of instructions and can be issued in parallel, and indicating at least another instruction of the plurality of instructions that can be issued after the first group of instructions;

issuing the first group of instructions in response to the instruction grouping information; and

coupling instructions in the first group of instructions to instruction pipelines appropriate for the instructions in the first group of instructions.

50. The method of claim 49 further comprises

after coupling instructions in the first group of instructions, issuing the at least another instruction in response to the instruction grouping information; and

coupling the at least another instruction to an instruction pipeline appropriate for the at least another instruction .

51. The method of claim 50 wherein the instruction grouping information also indicates which instructions of the plurality of instructions belong to a second group of instructions and can be issued in parallel after the at least another instruction .

52. The method of claim 49

wherein the instructions in the first group of instructions include instruction types; and

wherein coupling instructions in the first group of instructions further comprises determining the instruction pipelines appropriate for the instructions in the first group of instructions in response to the instruction types.

53. The method of claim 52 wherein the instruction types comprise opcodes.

54. The method of claim 50 wherein issuing the first group of instructions further comprises receiving the first group of instructions, the at least another instruction , and the instruction grouping information from the location in the memory.

55. The method of claim 50 wherein the first group of instructions comprises at least two instructions.

56. The method of claim 50 wherein the first group of instructions comprises at least one instruction.

57. The method of claim 55 wherein an instruction frame comprises the plurality of instructions and instruction grouping information, and

wherein the instruction frame includes at least the two instructions of the first group of instructions and the at least another instruction.

58. The method of claim 50 wherein the compiler explicitly identifies parallel executable instructions from the plurality of instructions.

59. The method of claim 52 wherein the instruction types comprise pipeline identifiers that identify the instruction pipelines appropriate for the instructions in the first group of instructions.

60. The method of claim 59 wherein the pipeline identifiers are determined by the compiler.

61. The method of claim 60 wherein coupling instructions in the first group of instructions comprises using a crossbar switch to couple the instructions in the first group of instructions to the instruction pipelines appropriate for the instructions in the first group of instructions in response to the pipeline identifiers.

62. The method of claim 50 wherein the memory is a cache and the location is a cache entry.

63. A computing system comprising:

a cache including a plurality of cache entries, a cache entry of the plurality of cache entries configured to store in parallel a plurality of software-scheduled instructions and instruction grouping information, the instruction grouping information configured to identify a first group of software-scheduled instructions from the plurality of software-scheduled instructions and to identify at least another software-scheduled instruction from the plurality of software-scheduled instructions, the at least another software-scheduled instruction to be issued after instructions in the first group of software-scheduled instructions.

64. The computing system of claim 63 wherein the instruction grouping information is also configured to identify a second group of software-scheduled instructions from the plurality of software-scheduled instructions, instructions in the second group of software-scheduled instructions to be issued after at least another software-scheduled instruction.

65. The computing system of claim 63

wherein the cache is also configured to issue the first group of software-scheduled instructions, the at least another software-scheduled instruction, and the instruction grouping information;

the computing system further comprising a group decoder coupled to the cache and configured to receive the first group of software-scheduled instructions, the at least another software-scheduled instruction, and the instruction grouping information, and to issue the first group of software-scheduled instructions in response to the instruction grouping information.

66. The computing system of claim 65 wherein the group decoder is also configured to issue the at least another software-scheduled instruction, after the first group of software-scheduled instructions in response to the instruction grouping information.

67. The computing system of claim 65

wherein each instruction in the first group of software-scheduled instructions includes an instruction type,

the computing system further comprising an instruction decoder coupled to the cache and configured to receive the instruction types of the instructions in the first group of software-scheduled instructions and to determine instruction pipelines appropriate for each of the instructions in the first group of software-scheduled instructions.

68. The computing system of claim 67 wherein the instruction type comprises an opcode, and the instruction decoder comprises an opcode decoder.

69. The computing system of claim 67 further comprising:

a pipeline coupler coupled to the group decoder and to the instruction decoder and configured to receive the first group of software-scheduled instructions and configured to couple each instruction in the first group of software-scheduled instructions to the instruction pipelines appropriate for the instructions in the first group of software-scheduled instructions.

70. The computing system of claim 69 wherein the pipeline coupler is a crossbar switch.

71. The computing system of claim 65 wherein the first group of software-scheduled instructions comprises at least two instructions.

72. The computing system of claim 71 wherein the first group of software-scheduled instructions comprises at least one instruction.

73. The computing system of claim 72 wherein an instruction frame comprises the plurality of software-scheduled instructions and instruction grouping information; and wherein the instruction frame includes at least the two instructions of the first group of software-schedule instructions and the at least another software-scheduled instruction.

74. The computing system of claim 63 wherein the cache is a superscaler cache.

75. The computing system of claim 63 wherein a compiler explicitly identifies parallel executable instructions from the plurality of software-scheduled instructions that form the first group of software-scheduled instructions.

76. The computing system of claim 67 wherein the instruction types comprise pipeline identifiers indicative of instruction pipelines appropriate for the instructions in the first group of software-scheduled instructions.

77. The method of claim 49 wherein the instruction grouping information also indicates instruction pipelines appropriate for the instructions in the first group of instructions; and

wherein coupling instructions in the first group of instructions to the instruction pipelines appropriate for the instructions in the first group of instructions is in response to the instruction grouping information.

78. The method of claim 77 wherein the instruction grouping information also indicates instruction pipelines appropriate for the instructions in the second group of instructions.

79. The computing system of claim 65 further comprising an switching unit coupled to the cache and configured to receive the instructions in the first group of

software-scheduled instructions and to couple the instructions in the first group of software-scheduled instructions to instruction pipelines appropriate for each of the instructions in the first group of software-scheduled instructions in response to the instruction grouping information.

80. The computing system of claim 79 wherein the switching unit is also configured to receive the at least another software-scheduled instruction and to couple the at least another software-scheduled instruction to an instruction pipeline appropriate the at least another software-scheduled instruction in response to the instruction grouping information.

81. (Amended) A computing system in which instructions are issued in parallel to processing pipelines, the computing system comprising:

a storage configured to store an instruction frame, the instruction frame including a plurality of instructions including a group of instructions and at least another instruction in parallel, instructions in the group of instructions to be issued in parallel, the at least another instruction to be issued at a time different from a time when the group of instructions is to be issued, the instruction frame also including data associated with the plurality of instructions, the data associated with the plurality of instructions indicative of which instructions in the plurality of instructions are included in the group of instructions, the data associated with the plurality of instructions also indicative of processing pipelines appropriate for the plurality of instructions, and the data associated with the plurality of instructions determined at a compile time; and

a switching circuit coupled to the storage, configured to issue the instructions in the group of instructions in parallel, to processing pipelines appropriate for the instructions in the group of instructions, in response to the data associated with the plurality of instructions.

82. The computing system of claim 81 wherein the group of instructions comprises at least two instructions.

83. The computing system of claim 81 wherein the group of instructions comprises one instruction.

84. The computing system of claim 81 wherein the switching circuit is also configured to issue the at least another instruction to a processing pipeline appropriate for the at least another instruction in response to the data associated with the plurality of instructions.

85. The computing system of claim 81 wherein the processing pipelines appropriate for the instructions in the group of instructions are respectively coupled to execution units appropriate for the instructions in the group of instructions.

86. The computing system of claim 85 wherein an execution unit appropriate for a first instruction in the group of instructions is a memory.

87. The computing system of claim 86 wherein an execution unit appropriate for a second instruction in the group of instructions is an arithmetic logic unit.

88. The computing system of claim 86 wherein an execution unit appropriate for a second instruction in the group of instructions is a floating point unit.

89. The computing system of claim 85 wherein an execution unit appropriate for one instruction in the group of instructions is a branch unit.

90. The computing system of claim 85 wherein a type of execution unit appropriate for a first instruction in the group of instructions and a type of execution unit appropriate for a second instruction in the one group of instructions are similar.

91. The computing system of claim 81 wherein the plurality of instructions in the instruction frame are determined at the compile time.

92. (Amended) A method for issuing groups of instructions in parallel to processing pipelines, the method comprising:

storing in a storage, an instruction frame, the instruction frame including a plurality of instructions including a group of instructions and at least another instruction in parallel, instructions in the group of instructions to be issued in parallel, the at least another instruction to be issued at a time different from a time when the group of instructions is to be issued, the instruction frame also including data associated with the plurality of instructions, the data associated with the plurality of instructions indicative of which instructions are included in the group of instructions, the data associated with the instructions also indicative of processing pipelines appropriate for the plurality of instructions, and the data associated with the plurality of instructions determined at compile time; and

issuing the instructions in the group of instructions in parallel to processing pipelines appropriate for the instructions in the group of instructions, in response to the data associated with the plurality of instructions.

93. The method of claim 92 further comprising:

during compile time, determining the plurality of instructions in the instruction frame.

94. The method of claim 92 wherein the group of instructions comprises at least two instructions.

95. The method of claim 94 further comprising, before issuing the group of instructions, issuing the at least another instruction to a processing pipeline appropriate for the at least another instruction in response to the data associated with the plurality of instructions.

96. The method of claim 92 wherein the processing pipelines appropriate for the instructions in the group of instructions are respectively coupled to execution units appropriate for the instructions in the group of instructions.

97. The method of claim 96 wherein a type of execution unit appropriate for a processing pipeline appropriate for a first instruction in the group of instructions is an arithmetic logic unit.

98. The method of claim 97 wherein a type of execution unit appropriate for a processing pipeline appropriate for a second instruction in the group of instructions is an arithmetic logic unit.

99. The method of claim 96 wherein a type of execution unit appropriate for a processing pipeline appropriate for a first instruction in the group of instructions is a floating point unit.

100. The method of claim 96 wherein a type of execution unit appropriate for a processing pipeline appropriate for a first instruction in the group of instructions is a memory unit and a type of execution unit appropriate for a processing pipeline appropriate for a second instruction in the group of instructions is a memory unit.

101. (Amended) A method of operating a microprocessor comprises:

compiling computer code to determine a frame of instructions;

storing in a memory storage the frame of instructions, the frame of instructions including a plurality of instructions and issue data, the plurality of instructions including at least a first instruction, a second instruction, and a third instruction, the issue data comprising data indicating that the first instruction is to be issued before the second instruction and the third instruction and that the second and third instructions are to be issued in parallel, and the issue data indicating respective processing units appropriate for the first instruction, the second instruction, and the third instruction; and

issuing the first instruction to a processing unit appropriate for the first instruction in response to the issue data; and

issuing the second instruction and the third instruction in parallel to respective processing units appropriate for the second instruction and the third instruction in response to the issue data.

102. The method of claim 101 wherein the processing unit appropriate for the first instruction is a memory unit.

103. The method of claim 102 wherein a processing unit appropriate for the second instruction is a memory unit.

104. The method of claim 102 wherein a processing unit appropriate for the second instruction is an arithmetic logic unit.

105. The method of claim 101 wherein issuing the first instruction to a processing unit comprises using a switching unit to couple the first instruction to the processing unit appropriate for the first instruction in response to the issue data.

106. (New) A computing system having a plurality of processing pipelines for executing groups of individual instructions, within very long instruction words, each individual instruction to be executed in each group being executed by different processing pipelines in parallel, the computing system comprising:

a main memory for storing a very long instruction word;

a very long instruction word storage, coupled to the main memory, for receiving the very long instruction word from the main memory, and for holding the very long instruction word, the very long instruction word including a predetermined number N of individual instructions, and including at least one group of M individual instructions to be executed in parallel, where M<=N, each individual instruction in the very long instruction word storage to be executed having an a pipeline identifier indicative of a processing pipeline for executing the individual instruction, and having a group identifier indicative of a group of individual instructions to which the individual instruction is assigned for execution in parallel;

a group decoder responsive to the group identifier for each individual instruction in the very long instruction word storage to be executed for enabling each individual instruction in the very long instruction word storage having a similar group identifier, to be executed in parallel by the plurality of processing pipelines; and

a pipeline decoder responsive to the pipeline identifier of each individual instructions in the very long instruction word storage to be executed for causing each individual instruction in a group of individual instructions having the similar group identifier to be supplied to the different processing pipelines.

107. (New) The computing system of claim 106 wherein M is greater than or equal to 1.

108. (New) The computing system of claim 106 wherein M is greater than 1.

109. (New) The computing system in claim 106, wherein the very long instruction word storage includes the at least one group of M individual instructions, and also includes group identifiers and pipeline identifiers for each individual instruction in the at least one group of M individual instructions.

110. (New) The computing system in claim 107, wherein each individual instruction in the at least one group of M individual instructions has associated therewith a different pipeline identifier.

111. (New) The computing system of claim 106, wherein the very long instruction word storage holds a first group of individual instructions to be executed in parallel and a second group of individual instructions to be executed in parallel after the first group, each individual instruction in the first group having associated therewith a first group identifier different from a second group identifier associated with each individual instruction in the second group, the first group and the second group being placed adjacent to each other in the very long instruction word storage.

112. (New) The computing system of claim 111 wherein:

the very long instruction word storage comprises a line in a cache memory having a fixed number of storage locations; and

the first group of individual instructions is placed at one end of the line in the cache memory, and the second group of individual instructions is placed next to the first group of individual instructions.

113. (New) A method of executing in a plurality of processing pipelines arbitrary numbers of instructions in a stream of instructions in parallel which have been compiled to determine which instructions can be executed in parallel, the method comprising:

in response to the compilation, assigning a common group identifier to a group of instructions which can be executed in parallel;

determining a processing pipeline for execution of each instruction in the group of instructions to be executed;

assigning a pipeline identifier to each instruction in the group;

associating the common group identifier and the pipeline identifier with the group of instructions;

forming a very long instruction word with a fixed number of the instructions including at least the group of instructions and the common group identifier as well as at least one other instruction having a different group identifier; and

storing the very long instruction word in a main memory.

114. (New) A method as in claim 113 further comprising:

placing the very long instruction word retrieved from the main memory into a very long instruction word register; and

executing the group of instructions in the plurality of processing pipelines in parallel.

115. (New) A method as in claim 114,

wherein the very long instruction word register holds at least two groups of instructions; and

wherein placing the instructions in the very long instruction word register comprises placing the group of instructions adjacent to the at least one other instruction having the different group identifier in the very long instruction word register.

116. (New) A method as in claim 115 wherein executing the group of instructions in parallel comprises:

coupling the very long instruction word register to a detection means to receive group identifiers associated with each instruction to be executed in the very long instruction word; and

supplying only instructions in the group of instructions to the processing pipelines in response to the group identifiers .

117. (New) In a computing system having a plurality of processing pipelines in which groups of individual instructions, within very long instruction words, are executable in parallel by processing pipelines, a method for supplying each individual instruction in a

group to be executed in parallel to corresponding appropriate processing pipelines, the method comprising:

retrieving a very long instruction word from a main memory;

storing in a very long instruction word storage the very long instruction word, the very long instruction word including groups of individual instructions to be executed in parallel, the groups of individual instruction to be executed in the very long instruction word having associated therewith pipeline identifiers indicative of the corresponding appropriate processing pipeline which will execute the instructions and group identifiers indicative of groups of instructions;

using the group identifiers in the very long instruction word to identify an execution group; and

using the pipeline identifiers to execute each individual instruction in the execution group in the corresponding appropriate processing pipelines.

118. (New) In a computing system having a plurality of processing pipelines in which groups of individual instructions, from a very long instruction word, are executable in parallel by the plurality of processing pipelines, an apparatus for routing each individual instruction in a particular group to be executed in parallel to an appropriate processing pipeline, the apparatus comprising:

a main memory for storing the very long instruction word;

a very long instruction word storage coupled to the main memory, for receiving the very long instruction word from the main memory and for holding the very long instruction word, the very long instruction word including groups of individual instructions, individual instructions to be executed in the very long instruction word storage having associated therewith pipeline identifiers indicative of processing pipelines for executing the individual instructions and also having associated therewith group identifiers to designate groups of individual instructions to which individual instructions are assigned, the pipeline identifiers and the group identifiers included in the very long instruction word;

a switching circuit having a first set of connectors coupled to the very long instruction word storage and a second set of connectors coupled to the plurality of processing pipelines; and

a router coupled to the very long instruction word storage and the switching circuit, responsive to the pipeline identifiers for routing each individual instruction in a group of individual instructions from connectors of the first set of connectors onto appropriate connectors of the second set of connectors, to thereby supply each individual instruction in the group of individual instructions to be executed in parallel to the appropriate processing pipeline.

119. (New) The apparatus of claim 118,

wherein the first set of connectors includes a set of first communication buses, one first communication bus for each individual instruction to be executed in the very long instruction word storage;

wherein the second set of connectors includes a set of second communication buses, one second communication bus for each processing pipeline; and

wherein the router comprises:

a set of decoders coupled to the very long instruction word storage, the decoders in the set for receiving as input signals the pipeline identifiers included in the very long instruction word storage and in response thereto for supplying as output signals switch control signals corresponding to each individual instruction in the very long instruction word storage; and

a set of switches coupled to the set of decoders and to the switching circuit, one switch of the set of switches at each intersection of each of the first set of communication buses with each of the second set of communication buses, each switch for receiving the switch control signals and for providing connections in response to receiving a corresponding switch control signal to thereby supply each individual instruction in the group to be executed in parallel to the appropriate processing pipeline.

120. (New) The apparatus of claim 119 further comprising:

a detection circuit coupled to the very long instruction word storage, for receiving the group identifiers included in the very long instruction word storage to be executed and in response thereto supply a group control signal; and

wherein the set of decoders are also coupled to the detection circuit for receiving the group control signal and in response thereto supply the switch control signal for only those individual instructions in the group to be supplied to the plurality of processing pipelines.

121. (New) The apparatus of claim 120,

wherein the detection circuit comprises a multiplexer coupled to receive the group identifiers included in the very long instruction word storage and in response thereto allow the group of individual instructions to be supplied to the plurality of processing pipelines.

122. (New) Apparatus as in claim 121 wherein the multiplexer supplies output signals to the set of decoders to indicate the group of individual instructions to be next supplied to the plurality of processing pipelines.

123. (New) In a computing system having a plurality of processing pipelines in which groups of individual instructions, within a very long instruction word, are executable by the plurality of processing pipelines, individual instructions in the very long instruction word to be executed having associated therewith group identifiers and pipeline identifiers, an apparatus for routing each individual instruction of a group of individual instructions to be executed in parallel to an appropriate processing pipeline of the plurality of processing pipelines, the apparatus comprising:

a main memory for storing the very long instruction word;

a very long instruction word storage coupled to the main memory, for receiving the very long instruction word from the main memory and for holding the very long instruction word the very long instruction word including groups of instructions to be executed in parallel, including pipeline identifiers and group identifiers;

a selection circuit coupled to the very long instruction word storage for receiving the group identifiers included in the very long instruction word, for determining in

response thereto a group of individual instructions to be executed in parallel, and for outputting a control signal;

a decoder circuit coupled to the selection circuit and to the very long instruction word storage, for receiving the control signal and the pipeline identifiers included in the very long instruction word, for determining in response thereto the appropriate processing pipeline for each individual instruction of the group, and for outputting switch control signals;

a switching circuit coupled to the decoder circuit , having a first set of connectors coupled to the very long instruction word storage for receiving the very long instruction word therefrom and a second set of connectors coupled to the plurality of processing pipelines, for coupling each individual instruction of the group to an appropriate processing pipeline in response to the switch control signals.

124. (New) The apparatus of claim 123,

wherein the first set of connectors comprises a set of first communication buses, one first communication bus for each individual instruction held in the very long instruction word storage;

wherein the second set of connectors comprises a set of second communication buses, one second communication bus for each processing pipeline;

wherein the decoder circuit comprises a set of decoders coupled to receive as first input signals the pipeline identifiers and as second input signals the pipeline identifiers; and

wherein the switching circuit comprises a set of switches, one switch for every intersection between each of the first set of connectors and each of the second set of connectors, each switch for providing connections, in response to receiving the switch control signals, between each individual instruction in the group to be executed in parallel to the appropriate processing pipeline.

125. (New) The apparatus of claim 124

wherein the selection means comprises a multiplexer coupled to receive the group identifiers for each individual instruction in the very long instruction word storage, and in response to the group identifiers, enable the decoder means to output switch control signals for each individual instructions of the group.

126. (New) The apparatus of claim 125,

wherein the multiplexer supplies a switch control signal to the decoder means to enable the decoder means to output switch control signals for each individual instruction of the group of individual instructions from the very long instruction word.

127. (New) In a computing system having a plurality of processing pipelines in which groups of individual instructions are executable, each individual instruction in a group executable in parallel by the plurality of processing pipelines, a method for transferring each individual instruction in a group to be executed through a switching unit having a first set of connectors coupled to a very long instruction word storage for receiving individual instructions therefrom, a second set of connectors coupled to the plurality of processing pipelines, and switches between the first set and the second set of connectors, the method comprising:

retrieving the very long instruction word from a main memory;

storing in the very long instruction word storage, the very long instruction word, the very long instruction word having a set of individual instructions including at least one group of individual instructions to be executed in parallel, individual instructions in the at least one group having associated therewith pipeline identifiers indicative of the processing pipeline which will execute that individual instruction, the very long instruction word storage also including at least one other individual instruction not in the at least one group of individual instructions, the at least one other individual instruction also having associated therewith the pipeline identifiers; and

using the pipeline identifiers to control the switches between the first set of connectors and the second set of connectors to thereby supply each individual instruction in the at least one group to be executed in parallel to an appropriate processing pipeline.

128. (New) A method as in claim 127 wherein the step of using the pipeline identifiers comprises:

supplying the pipeline identifiers to individual decoders of a set of decoders, each decoder of which provides an output signal; and

using the output signals of the sets of decoders to control the switches between the first set of connectors and the second set of connectors to thereby supply each individual instruction in the at least one group to be executed in parallel to an appropriate processing pipeline.

129. (New) A method as in claim 128 wherein individual instructions in the storage further includes group identifiers associated therewith to designate among the instructions present in the very long instruction word storage, which of the individual instructions may be simultaneously supplied to the plurality of processing pipelines, and the method further comprises:

supplying a group of instructions to be executed by the processing pipelines together with the group identifiers to a selector;

using the group identifiers to provide output determination signals; and

using both the output determination signals and the output signals to control the switches between the first set of connectors and the second set of connectors to thereby supply each instruction in the at least one group to be executed in parallel to the appropriate processing pipeline.

130. (New) In a computing system having a plurality of processing pipelines in which groups of individual instructions are executable by the plurality of processing pipelines, a method for supplying each individual instruction in a group of individual instructions to be executed in parallel to an appropriate processing pipeline, the method comprising:

retrieving a very long instruction word from a main memory;

storing in a very long instruction word storage the very long instruction word retrieved from the main memory, the very long instruction word including groups of individual instructions to be executed in parallel, the individual instructions having associated therewith pipeline identifiers indicative of processing pipelines which will execute the individual

instructions and having associated therewith group identifiers indicative of a group identification;

comparing the group identifiers to an execution group identifier of those instructions to be next executed in parallel; and

using the pipeline identifiers to control switches in a switch having a first set of connectors coupled to the very long instruction word storage for receiving the very long instruction word therefrom and a second set of connectors coupled to the plurality of processing pipelines to thereby supply each individual instruction in the at least one group to be executed in parallel to the appropriate processing pipeline.